

Kalman Filter Tracking on Parallel Architectures

Giuseppe Cerati¹, Peter Elmer², Steven Lantz³, Kevin McDermott⁴,
Dan Riley⁴, Matevž Tadel¹, Peter Wittich⁴, Frank Würthwein¹, Avi
Yagil¹

¹ University of California - San Diego , La Jolla, CA, 92093, USA

² Department of Physics, Princeton University, Princeton, NJ 08540, USA

³ Center for Advanced Computing, Cornell University, Ithaca NY 14853, USA

⁴ Laboratory of Elementary Particle Physics, Cornell University, Ithaca NY 14853, USA

E-mail: cerati@cern.ch

Abstract. Power density constraints are limiting the performance improvements of modern CPUs. To address this we have seen the introduction of lower-power, multi-core processors, but the future will be even more exciting. In order to stay within the power density limits but still obtain Moore's Law performance/price gains, it will be necessary to parallelize algorithms to exploit larger numbers of lightweight cores and specialized functions like large vector units. Example technologies today include Intel's Xeon Phi and GPGPUs.

Track finding and fitting is one of the most computationally challenging problems for event reconstruction in particle physics. At the High Luminosity LHC, for example, this will be by far the dominant problem. The need for greater parallelism has driven investigations of very different track finding techniques including Cellular Automata or returning to Hough Transform. The most common track finding techniques in use today are however those based on the Kalman Filter [2]. Significant experience has been accumulated with these techniques on real tracking detector systems, both in the trigger and offline. They are known to provide high physics performance, are robust and are exactly those being used today for the design of the tracking system for HL-LHC.

Our previous investigations showed that, using optimized data structures, track fitting with Kalman Filter can achieve large speedup both with Intel Xeon and Xeon Phi. We report here our further progress towards an end-to-end track reconstruction algorithm fully exploiting vectorization and parallelization techniques in a realistic simulation setup.

1. Introduction

Pile-up (PU) represents a challenge for HEP event reconstruction, both in terms of physics performance and in terms of processing time. In fact, for PU values exceeding 100, as expected at the High Luminosity LHC (HL-LHC), the time needed for event reconstruction diverges (Fig. 1); due to power density limitations to Moore's law, such a large increase is not sufficiently compensated by an increase in CPU clock frequency. For this reason, the reconstruction model currently used in most HEP experiments, based on traditional computers both for offline and online processing, cannot be sustained at HL-LHC without compromises between timing and physics performance that will impact the final sensitivity of the experiments.

As a solution to this problem we investigate a transition to highly parallel computing architectures such as e.g. Intel's Xeon Phi and NVIDIA GPGPUs. The challenge in this context is that, due to distinct features of these architectures, a simple porting of the current, 'serial',

implementation of the algorithms would be highly suboptimal: code and data-structures need to be redesigned with a strong emphasis on hardware capabilities and limitations.

We initially focus our study on the Xeon Phi architecture. The main reason for this choice is that it shares many features with more conventional architectures, like the Intel Xeon, so that we can optimize and test the algorithm performance on both architectures at the same time; however, there is no real prejudice on the choice of the architectures and in the future we plan to explore also other options, including GPGPUs. Xeon Phi coprocessors are characterized by 60 cores running up to 4 threads each and featuring 512-bit wide vector units; thus, in order to achieve optimal performance all cores must be kept occupied, usage of vector units must be maximized, and code branching must be minimized.

We tackle the problem starting from the the most challenging algorithm, track reconstruction; tracking is by far the most time consuming process in event reconstruction (Fig. 1) so that, if the proposed approach does not work in this case, there is little to gain from the rest of event reconstruction. Other algorithms, such as vertexing, jet clustering and so on will be studied at a later stage. The present report describes the result of early studies and outlines the next steps of the project.




Figure 1. Timing vs luminosity for $t\bar{t}$ +PU samples with 25 ns bunch crossing. Samples with different average pile-up are used as reported on the plot. Shown is timing of full CMS reconstruction and of the iterative tracking sequence [1].

2. Setup and Previous Results

Below is a brief summary of the test setup and of previous results; a more detailed description can be found in [3].

In order to study the problem with maximum flexibility and with gradually increasing complexity, we developed a standalone tracking code based on Kalman Filter and running on a simplified setup consisting of an ideal barrel geometry, a uniform and longitudinal magnetic field, gaussian-smeared hit positions, a particle gun simulation with flat transverse momentum distribution between 0.5 and 10 GeV, no material interaction and no correlation between particles nor decays.

The track reconstruction process can be divided in 3 steps: track seeding, building and fitting. The track fit is the simple application of the Kalman Filter to a pre-determined set of hits, so it was a natural choice as a starting point. It is implemented using Matrimplex, a library for vectorized matrix operations we developed. Matrimplex is a matrix-major representation, where vector units elements are separately filled by the same element from different matrices; in other

terms, n matrices - and thus n tracks - work in sync so that vectorization is used as an additional resource for parallelism. The track fit using Matriplex gives same physics results and, even in serial case, it is faster than an equivalent version based on SMatrix.

We tested the track fit both on Intel Xeon and Xeon Phi (native application) with OpenMP, obtaining similar qualitative results. We observed a large speedup both from vectorization (4x/8x on Xeon/Xeon Phi) and parallelization (10x/80x). The effective utilization efficiency of vector units is about 50%; time scaling is close to ideal in case of 1 thread/core, while there is some overhead with 2 threads/core. Main performance limitations are related to data availability in L1 cache and data re-packing in Matriplex format.

As steps forward, we followed two distinct lines of development. On the one side we consolidated the track fitting results, identifying the performance bottlenecks to vectorization and progressing towards a fully realistic setup. On the other side, we tackled the most time consuming part of the track reconstruction algorithm, redesigning the track building in order to exploit vectorization and parallelization.

3. Improvements to Track Fit

A detailed profiling of the code showed that the relative fraction of time used for data input is large with respect to total fit time, where input refers both to data transfer to L1 cache and re-packing into Matriplex format. Figure 2 compares the time spent for data input using different methods to the actual fit time with Xeon Phi on a single-threaded execution using the vector units to the extent possible. We explored different approaches (methods 1-3) and different intrinsics (4-6), either with a direct copy approach (1 and 6) or with a two-stage copy using temporaries (2-5), resulting in substantially different performance. The best method (#4) relies on two-stage copy and employs the vgather intrinsic.




Figure 2.

1. scatter data into Matriplexes, track by track, using simple for-loops
2. copy 16 tracks into contiguous memory, transpose into Matriplexes via loops
3. copy 16 tracks into memory, use Intel MKL to transpose each Mplex
4. copy 16 tracks into memory, vgather into Matriplexes by rows
5. copy 16 tracks into memory by rows, vscatter into Matriplexes
6. scatter data into Matriplexes, track by track, using Intel intrinsics

We also made significant progress towards the larger goal of implementing an end-to-end tracking chain in a realistic simulation, with track fitting as the final step. Our initially simple setup was designed to allow increases in complexity, so we can add more realistic features in an incremental way. In particular, we included a more realistic detector description implementing different barrel-like layouts by using the USolids geometry package [4]. Such changes are transparent to the actual fitting code since the track parameters are propagated to the radius of the next hit in global coordinates. We verified that the timing and physics performance of the track fit are not affected (Fig. 3 and 4). Further improvements include the options to account

for the effects of multiple scattering in the simulation and, instead of fitting the track hits as obtained from the simulation, to fit the tracks returned by the track building.




Figure 3. Pull of the track transverse momentum using an barrel geometry composed of cylindrical layers.




Figure 4. Pull of the track transverse momentum using an barrel geometry composed of polyhedral layers.

4. Track Building

The track building is the part of the track reconstruction process that, starting from a proto-track made of a 2-3 hits (the seed), identifies all other hits belonging to the track. From the computing point of view, it is by far to most expensive task in HEP event reconstruction.

While it entails similar calculations as the track fitting (i.e. track parameter propagation and update, χ^2 of each measurement), it adds two big complications to the problem. First, while during fitting the set of hits to process is already defined, in the building case a large number of hits is present at each layer, typically of the same order of the track multiplicity in the event. Second, when more than one compatible hit is identified on the same layer, the combinatorial nature of the algorithm requires that a new track candidate is created for each compatible hit.

Data locality is the key for reducing the large number of hits problem, so that only a reduced set of hits is considered at a given time and thus needs to be available in lower caches levels. For this purpose, we partition the space both in pseudorapidity (η) and in azimuthal angle (ϕ). We exploit the fact that tracks do not bend in η to define self-consistent η partitions which are redundant in terms of hits, so that there are no boundary effects and track candidates never search outside their η bin (Fig. 5); η bins constitute a natural splitting for thread definitions. A simple partitioning in ϕ is used for fast look-up of hits in the 3σ compatibility window, where σ is the uncertainty on the propagated track ϕ on the considered layer.




Figure 5. Diagrams illustrating one single η bin (left) and the full η partitioning (right).

The two issues described above (i.e. hit multiplicity and branching of candidates) can be addressed separately by factorizing the problem in two stages. In a first version of the algorithm,

at each layer all hits within the compatibility window of a given track are probed and only the hit with the lowest χ^2 is chosen, so that there is no multiplication of candidates; this simple and fast version already achieves good hit collection performance: running over 10 events with 20k tracks each, 70% (93%) of tracks are reconstructed with $\geq 90\%$ (60%) of the hits. In this setup we studied the performance in terms of vectorization, obtaining a maximum speedup of $>2x$ both on Xeon and Xeon Phi (Fig. 6 and 7). The scaling with vector unit size is monotonic on Xeon while an overhead is observed when vectorization is enabled on Xeon Phi, followed by a gradual speedup; the usage of prefetching and gathering intrinsics gives a further time reduction when the vector units are maximally used. It should be noted that, with respect to the fitting case, worse vectorization performance is expected since the latency due to data input is amplified by the much larger number of hits processed at each layer.




Figure 6. Track building time (best hit version) vs number of elements in vector unit for Xeon processor.




Figure 7. Track building time (best hit version) vs number of elements in vector unit for Xeon Phi.

The second version is more advanced and targets ultimate physics performance. In this case we allow the branching of candidates, so that at each layer new candidates are created for all hits passing a χ^2 cut of 30; also, in order to account for outlier hits and detector inefficiencies, candidates are allowed to miss the hit on one layer. Up to 10 candidates per seed are considered and, when this number is exceeded, candidates are sorted based on the number of hits and the total χ^2 and those in excess are discarded. Running over the same set of events of 20k tracks each, the hit collection performance is significantly improved and 85% (95%) of the tracks are reconstructed with $\geq 90\%$ (60%) of the hits. We tested the performance of this algorithm in terms of parallelization. Threads are evenly distributed across 21 eta bins, so that when the number of eta bins (nEtaBins) is a multiple of the number of threads (nThreads), $n = n\text{EtaBins}/n\text{Threads}$ eta bins are processed in each thread, while for nThreads multiple of nEtaBins, seeds in each eta bin are distributed across $n = n\text{Threads}/n\text{EtaBins}$ threads. A speedup of 5x on Xeon and $>10x$ on Xeon Phi is achieved (Fig. 8 and 9); on Xeon Phi, we observe a saturation above nThreads=42 that will require further investigations.

5. Conclusions and Outlook

In summary, a first implementation of a vectorized and parallelized track building in a simplified setup achieves significant speedup both on Xeon and Xeon Phi: 2x from vectorization on both architectures, 5x on Xeon and 10x on Xeon Phi from parallelization. Track fitting studies have been consolidated by implementing an improved data input method and a more realistic detector geometry description.

The current results are definitely promising but many developments are still needed to deploy a complete tracking on parallel architectures. The comparison with the scaling for




Figure 8. Track building time vs number of threads for Xeon processor.

